Pulsar IIa Design Review Summary

with answers to reviewer's questions

Summary of changes to the Pulsar IIa board:

- Added fiducial markers.
- Testpoints added to some of the unused FPGA I/O pins.
- FMC output clocks F1_CLK0 and F1_CLK1 have been routed to MRCC clock inputs in Bank 13. These clocks may then be distributed to neighboring banks 12 and 14 using regional clock resources.
- FMC output clocks F2_CLK0 and F2_CLK1 have been routed to MRCC clock inputs in Bank 16. These clocks may then be distributed to neighboring banks 15 and 17 using regional clock resources.
- Soldered-in fuses changed to socketed OMNI-BLOK fuses and moved from side 2 ("solder side") to side 1 ("component side"). Where these surface mount fuses connect to internal planes multiple large vias are used.
- In order to support PICMG 3.8 hot swap functionality the RTM should appear like an AMC mezzanine card to the front board. The Pulsar IIa board presented for review did not properly implement hot swap, so the following changes were made:
    - Added I2C bus isolator chip.
    - Added high side power switches to control the +3.3V and +12V power to the RTM.
    - Added logic to monitor the PS# line and drive the ENABLE# line. The IPMC microcontroller now has the ability to reset the MMC on the RTM.
    - The IPMC microcontroller has the ability to monitor the +12V RTM current.
- The Base Interface 100BASE-T transformer filter capacitor (C5) has been increased to 1000pF 2kV. Added a 75 ohm resistor in parallel with C5.

Summary of changes to the RTM board:

- Removed the opto-coupler on the ENABLE# line; now the DC-DC converter is configured to turn on as soon as 12V power is applied to the board. This is proper "AMC" hot swap behavior; the 12V power is now controlled by the front board.
- Removed the I2C bus isolator; this isolator is now on the front board.
- Added a logic gate so that the ENABLE# signal properly resets the MMC microcontroller.
- The soldered-in fuse was replaced with a socketed OMNI-BLOK fuse. Multiple large vias are used to connect this surface mount fuse to the power planes.

The updated designs are available on our website:

http://www-ppd.fnal.gov/EEDOffice-w/Projects/ATCA/

Ted's Email invitation:

From: Tiehui Liu <thliu@fnal.gov>
To: Tiehui Liu <thliu@fnal.gov>
CC: Theresa M Shaw <tshaw@fnal.gov>, Alan G Prosser <aprosser@fnal.gov>,
Ryan A Rivera <rrivera@fnal.gov>, Neal G Wilcer <wilcer@fnal.gov>,
Jeffry L Andresen <andresen@fnal.gov>, Lorenzo Uplegger <uplegger@fnal.gov>,
"John T. Anderson x8885 08532N" <jta@anl.gov>, "Drake, Gary R." <drake@anl.gov>,
Mircea Bogdan <bogdan@edg.uchicago.edu>, Fukun Tang <fukun.tang@gmail.com>,
"Xu, Hao" <haoxu@bnl.gov>, Peter J Wilson <pjw@fnal.gov>,
Erik E Gottschalk <erik@fnal.gov>, Simon Kwan <swalk@fnal.gov>, Marcus H Larwill
<larwill@fnal.gov>, Jamieson T Olsen <jamieson@fnal.gov>,
Yasuyuki Okumura <okumura@fnal.gov>
Subject: Re: review of Pulsar IIa ATCA based prototype board design (Nov.
14th, 1:30pm to 4:30pm)
Date: Thu, 15 Nov 2012 11:21:53 -0600

Dear All,

We would like to invite you to review our Pulsar IIa ATCA based Data Formatter *prototype* design. **The review will be held on Nov. 14th, 1:30pm to 4:30 pm**, at Wilson Hall, 10th floor, Room 1000 (next to men's room).

The review materials can be found here:

http://www-ppd.fnal.gov/EEDOffice-w/Projects/atca/

The focus of this (informal) prototype design review will be on the details of technical implementations, such as meeting FPGA/components guidelines, high speed serial lines implementation, power distributions, mechanical, general operations such as downloading firmware, interface implementation between mezzanine and motherboards ... etc, to name a few.

In particular, since this is the very first ATCA prototype from Fermilab, one focus will be on meeting ATCA requirements.

Many thanks in advance,

Ted (with Jamieson and Yasu)

Answers to reviewer's questions

Comments from John Anderson (ANL):

Few comments here, although I do wish to say that it was a distinct pleasure to participate in such a lively, yet friendly, review. Overall the board design is very impressive and it was a most excellent presentation. My comments are as follows.

1) I am not absolutely certain that you need a transformer to interface to the backplane. Might be worth double-checking whether AC coupling is sufficient.

PICMG section 6.5.4. suggests that the Base Interface should use a transformer to meet 10/100/1000BASE-T isolation requirements.

2) Agree that fuses should be in sockets if you can fit them, and also that fuses in general should have multiple parallel vias. While we didn't look in detail, would also suggest you insure a sufficient number/type of vias in parallel for your big power converter. FNAL safety reviews love to look at the current density in vias.

Littelfuse SMT fuses have been added to side 1. We are using three 0.5mm vias per pad when the fuse connects to the plane. The larger switching power converters are through hole parts.

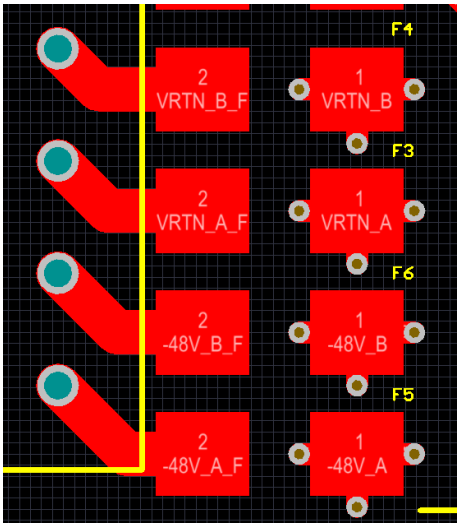

3) Suggest that adhesive heat sinks will probably suffice for your FPGAs.

Clip on or adhesive heat sinks will be used:

4) Please insure that the machine shop you select for front panel work can handle C-channel blanks and that they won't mangle the ESD gasket of blank panel kits.

Front panel blanks are on hand; Johnny Green to check on local machine shop capabilities.

5) A simple but easily overlooked item is that ATCA handles come in two kinds - push on the switch directly and "lean" on the switch as the tang of the handle passes by. Insure the orientation of the handle switch matches the handles you've bought.

We are using a “lever style” handles and matching Southco handles.

6) Given your layer stack-up I would guess that back-drilling your high speed vias may be a cost-effective form of insurance, and I recommend you give it due consideration if speed truly warrants. A potentially useful reference article is at http://onlinelibrary.wiley.com/doi/10.1002/ecjb.20152/pdf and you may also wish to look at the various design articles found at ultracad.com. Generally I think you are wise to be planning on the use of Rogers material if your individual transmission line pairs are truly going to run at rates in excess of 3Gbps, but if further system analysis shows that you only need an aggregate of 10Gbps per *channel* - four pairs - then FR4 is almost certain to suffice. The loss tangent for FR-4 does not generally become problematic until about 3GHz. Overall my feeling is that the Rogers vs. FR-4 choice should not be agonized over if the cost is relatively equivalent; Rogers materials are common and well understood by most competent board houses.

Our prototype boards will use RO4350 cores and pre-preg material and we will evaluate high speed performance both on the RTM and ATCA full mesh backplane. We will consider back-drilling high speed vias.

Comments From Hao Xu (BNL):

It is a beautiful design and really ATCA compliant. The layout is very clean and follows the high speed signals design rules. I'm sure you will be experts in ATCA soon when you get the board and play with it.

1. ATCA standard, 1.1 Channels 1~9 are selected to connect to the full-mesh backplane in current design. Just remind that only slots 1~10 well be full-mesh connectivity by this prototype.

Correct. The prototype boards will only reside in the first 10 slots of the full mesh backplane.

1.2 PICMG 3.0 supports hot swappable or non hot swappable RTM. Since the management power (+3.3V) is provided directly from front board, RTM will not support hot swap. Management circuits should be added on front board if hot swap needed in future. Management power/payload power should be controlled by IPMC on front board, SCL/SDA buffer should be moved to front board.

This is not described in PICMG 3.0/3.8, but we can get an example from MTCA.4, as shown in Fig. 1. Another example block diagram of the basic RTM and front board management circuitry is shown in Fig. 2. There are many chips can handle hot swap well. LTC4223-1, which is a Linear product I used on Compute Node V3, is a good choice.

Circuitry has been added to make the RTM operate just like a hot-swappable AMC mezzanine card:

1) An isolator chip has been added to the I2C lines going to the RTM; this isolator is controlled by the IPMC.
2) The IPMC independently controls the +3.3V and +12V power supplied to the RTM. The IPMC now monitors the current on the 12V power rail.
3) A NAND gate has been added to properly drive the ENABLE# line to reset the RTM MMC.

1.3 ATCA RTM Power/Cooling. Please be advised to check with the power consumption of the RTM since the power is limited to 25W by PICMG 3.8 R5.1. Cooling evaluation will be needed if the RTM power over 25 W.

Our RTM power calculations are based on 1.5W per QSFP+ (Avago AFBR-79E4Z) and 1W per SFP+ (Avago AFBR-703SDDZ). If all transceivers are used and running 10Gbps we calculate 18W per RTM worst case.

1.4 IPMI/IPMC. IPMI/IPMC is not trivial. Too many works have to be done to meet the specification. However, many functions/commands will not be useful for DF system at all. Starting from basic functions will be better to save time and resources.

Initially the IPMI functionality can be omitted to bring up the board. However we will be working closely with other Engineers to implement the minimum IPMI feature set. Our efforts will be open source and freely available.

1.5 Fuse holder. OMNI-BLOK from Littelfuse may be quite suitable for ATCA board. I used it on Compute Node without any problem.

We have switched to Littelfuse OMNI-BLOK fuses and moved them from side 2 to side 1.

2. GTX transceivers. It is recommended to use AC coupling capacitors for connection to transmitter and receiver. As a universal module, it would be better to put AC coupling capacitors on channels connected to RTM. So it's not needed to check with what kind of transceivers used on RTM every time. Try to do BERT between different boards. Use PRBS 31-bit instead of PRBS 7-bit to get more accurate result.

AC coupling capacitors are used on the GTX receivers to the Fabric Interface, per PICMG specification. For RTM optics we plan to use transceivers with built-in AC coupling capacitors. We will consider modifying the RTM software to check for the presence of "non-standard" QSFP and SFP transceivers installed on the RTM.

3. DDR3 controller firmware design. Problems were found when people tuned DDR3 with some Xilinx FPGA at high performance. Kintex7-325T is involved in. Use the latest ISE software to generate the controller with new flight time data to avoid this kind of problem. Details is listed on Xilinx website, http://www.xilinx.com/support/answers/51296.htm

We are using the latest Xilinx 14.3 tools and have verified the DDR3 pinout using Coregen.

4. Power on/off power supply sequencing. As mentioned in Xilinx K7 datasheet, the recommended power-on sequence is VCCINT, VCCBRAM, VCCAUX, VCCAUX_IO, and VCCO. There is a similar request for power supply of GTX transceivers too. Please make sure the design meets the requirements.

All switching regulators are powered from the +12V rail. A few low power rails use LDO linear regulators and these track the switching regulator outputs. Using typical values from the datasheets shows that all switching and LDO linear regulators start ramping up at ~3ms after the 12V power starts ramping up. All regulators reach their final voltage ~6ms after the 12V power starts ramping up.

Xilinx indicates that the parameter TVCCO2VCCAUX should be less than 800ms and we are less than 1ms for all VCCO rails.

We have consulted with our Xilinx Field Application Engineer and he indicates that simultaneous monotonic power sequencing is fine for the 7-series production parts.

5. FR4 or Rogers. In my experience, there isn't any different between FR4 and Rogers when the line rate is up to 8Gbps. It is not very clear when the board runs at 10Gbps since there are one bad and two good channels on my FR4 board. However, it will be safe to use Rogers for a board with so many high speed signals. Please let me know your test results.

We will use Rogers for the first prototype boards and evaluate cost and performance versus using FR4 on future designs.

From Mircea Bogdan (UC E-Shop):

The option of installing FPGA heat sinks was mentioned during the review. Are there any holes for the screws? I can't see them in the artwork.

We are planning on using a clip on passive heat sink, similar to the device used on the Kintex development board (without the fan). Another option is to use a self-adhesive heat sink.

From Fukun Tang (UC E-Shop):

I am very impressed with the beautiful work completed by Fermilab team. As prototype run, I think the schematic design and layout are completed to meet the goal. I am looking forward to the test results of the demonstrator in a few months. Following are my comments and suggestions based on my best understanding to the information you provided both in your design files and presentations in the design review meeting. Some of them may be quickly added into prototype, some of them can be considered later in your production run. It is all your choices.

Schematics Design

A. Sheet 3/17: The USB_VBUS, USB_D+ , USB_D- can be bidirectional clamped by diodes between 3.3V and ground, not just clamping from ground only.

We are following Silicon Labs CP2102 recommendations, see datasheet Figure 5.

B. Sheet 4/17: As we talked on the meeting, we should add a 75 ohm resistor to ground from C5. Commercially, a 75 ohm resistor should connect from C5 to RJ45 shielding. Since you don't use RJ45, it should connect it to ground instead. The 75-ohm resistor to ground will provide a path to reduce the unbalanced differential current that will turn to a low frequency noise after filtered by C5. In addition, it also helps provide the impedance matches, i.e. ~110 ohms seen from any center pins of the transformers. Well terminated center pins will help reduce both common mode and differential mode EMI noise.

Added 75 ohm resistor from net ETH_FILT to ground.

C. C5 should be rated at above 2KV, it will help prevent from lighting damage induced from Ethernet cable. It is all about the safety.

Changed C5 to 1000pF 2000V capacitor.

D. Please make sure the U3 can receive signals in a common mode voltage of 3.3V, as shown in your schematics.

Yes, the board design matches the Ethernet PHY chip example design (DP83848C, Figure 11).

E. Sheet8/17: I hope you will not change fuses too frequently since you need remove and mount fuses by solder, it has a risk causing the pads off the board. I suggest to use SMD fuse blocks if the height fit the space. See the link for low profile SMF blocks, so that you don't need solder fuse any more.

Fuses have been replaced with Littelfuse OMNI-BLOK 154 series fuses with surface mount holders. The fuses have also been moved to side 1 for easy access. Fuse ratings have been changed to accurately reflect the maximum board power usage (240W, 5A @ 48V). Fuses F3 and F4 change from 10A to 8A, or 60% greater than the expected max current of 5A. The power return fuses F5 and F6 change from 12A to 10A. (See PICMG 3.0 section 4.1.4 for fusing example.)

The -48V supply and return traces have been replaced with internal split planes and the fuse SMT pads use 3 x 0.5mm vias. We verified that the SMT fuses have a DC rating of 125V.

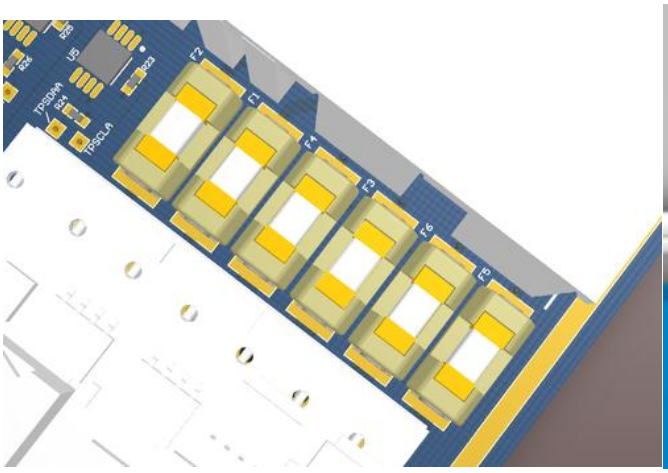

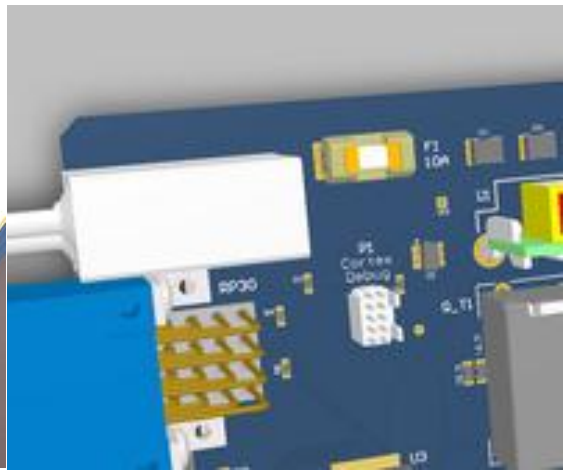

The power input fuse on the RTM has also be replaced with an OMNI-BLOK fuse holder.

F. Sheet15/17: The LDO U12 will take 0.8A from 3.3V to 1.8V, if you take an input voltage from 2.5V, then you can reduce the power consumption of 0.64 watts per LDO. The DF card is a power hungry card, if you can reduce power in any spots you know, it counts.

Xilinx power estimator indicates that the MGTVCCAUX uses 41mA, worst case. Power dissipation in U12 is negligible.

G. The layout is very beautiful, it complies with the high speed signal integrity rules.

Thanks!

H. I fully agree with you using Rogers for this board in prototype before you get anything out. The FR4 may work if the transceivers have pre-emphasis and equalizer features available.

We plan to start with Rogers material and if possible, may consider scaling back to FR4 in the future.

I. General “blind” suggestions: Is it possible to add some of test points on board for FPGA test? I know you can use signal taps to debug your firmware, however, if you have a few key test points laid out for probing, it will be very useful too. However, none should be on any high speed lines since it will degrade the signal integrity.

Added some general purpose FPGA test points on unused I/O pins.

J. Dumb question: I like the LV power controller. However, if the microprocessor goes wild or the board is running into over current, what will happen to the controller and FPGAs? What is the failure mode?

We plan to use the watchdog timer to insure that the microcontroller remains in a normal operating state. All switching regulators have over-voltage, over-current, and over-temperature protection. Furthermore the microcontroller can directly monitor board power consumption through the bus converter I2C interface. If the bus converter is running too high the microcontroller shall inform the shelf manager and shutdown the bus converter.

Firmware and Test

A. You definitely have lots of bandwidth available based on the trigger conditions if you push the transmission rate up to 10Gbps per link, I suggest put some error correction schemes (such as Hamming code) for the communications between DF and AUX. I believe there will be lots of work to explore in firmware design, including the feasible, intelligent methods both for self-test and normal operations.

We plan to use SLINK protocol for transmissions into and out of the DF system. Aurora protocol will be used for all high speed communication between FPGAs within the system.

B. The self-looped testing capability is wonderful.

We have many options for looping back in this architecture and we intend to use all of them!

Comments From Jeff Andresen (Fermilab):

I have the following comments about the Pulsar IIa review. I thought the overview presentation was helpful. This is an exciting project. The PCB layout is excellent. The performance of the design should not be limited by the PCB layout. The optimization of the FPGA pins contributed greatly to how good the PCB layout is. Designing the prototype for 10 Gbps is a good idea for gaining knowledge for this project and future projects. Using Rogers PCB material is important for high speed designs and is not a major cost increase compared to the total project cost. It would be unwise to risk the success of the design to save some money by using FR4.

From my experience, the biggest degradation of signal quality will come from the FMC and rear connectors. We have tested a design using the FMC connector at 10 Gbps with good success. However, when looking at the design with a TDR, there is a major impedance disruption going through the FMC connector. In spite of this, the design works. Vias were barely noticeable compared to the FMC connector. With the FMC connector being surface mount, I would think it will be better than the press fit through hole Zone 3 rear connectors. We have not yet tested a design with those connectors. I know people like the idea of have the 10 Gbps devices on the RTM. It will be good to find out how well it works with the longer trace lengths and the connector that using the RTM requires. There is little question that having the 10 Gbps device near the FPGA with the fiber connector on the front panel would have the best performance. Ideally, you would have the 10 Gbps optical device on the FPGA.

In our design the FMC signals will likely not exceed 1Gbs LVDS. Everything I can find indicates that the press fit backplane connectors ("Advanced Differential Fabric") are rated for 10Gbps operation. Our ELMA backplane is also specified for 10Gbps per port or 40Gbps per channel. We will see!

Having a 10 Gbps device directly on the Pulsar IIa PCB would verify how well the design operates without the connectors. I do not remember if there is a 10 Gbps directly on the Pulsar IIa PCB. A PCB design operating at 10 Gbps is very challenging. The thought of it operating with a low bit error rate through connectors is extremely challenging. It would be good to TDR the completed designs.

Both of you are to be commended for an excellent design.

Thanks!!!

Comments from Ryan Rivera (Fermilab):

Just wanted to quickly say that I was very impressed by everything presented in the review today. It's obvious that a lot of thinking and planning has gone into the project. I look forward to seeing the board in action!

Thanks!!!

And Thanks to all for your participation and feedback on this design!